UCRL-JC-125001 Abs

Experiences with a massively parallel coupled oceanic-atmospheric general circulation mode l.*

Michael Wehner
Climate Systems Modeling Group
L-256
Lawrence Livermore National Laboratory
Livermore, CA 94551
mwehner@llnl.gov

The Climate Systems Modeling Group at Lawrence Livermore National Laboratory (LLNL) has recently constructed a coupled oceanic-atmospheric general circulation model (O AGCM) specifically targeted for massively parallel processing (MPP) computers. Since MPP architectures are rapidly evolving, portability has been a design standard since the code's inception. However, since the promise of increased computational speed is the reason for going to massively parallel machines in the first place, we have attempted to achieve portability without adversely affecting performance. To this end, we have combined domain decomposition of the individual model components with a task decomposition scheduling algorithm to further parallelism and scalability.

Portability issues are mainly solved by the usage of the MICA (Macro Interface for Communication and Allocation) toolset. As the name implies, we use this set of M4 macros to implement dynamic memory management and message passing syntaxes where there are no FORTRAN 77 standards. (M4 is a widely available UNIX preprocessor.) We have successfully ported our code, with these macros, to a wide variety of MPP machines. This includes the Cray T3D, IBM SP2, Meiko CS2, Intel Paragon, and TMC CM5 as well as other uniprocessing vector and workstation architectures. With a wide variety of message passing libraries included in this toolset, we are free to choose the most favorable one to optimize performance on a particular machine.

A combination of domain and task decomposition is the centerpiece of the parallel coupled OAGCM. The atmospheric and oceanic components of the model are each allowed to have independent two dimensional domain decompositions following Richardson's (1922) idea. These independent sections of the code are simultaneously executed on separate sets of processors. A flux coupler module is then executed by both processor sets at the designated coupling time. This coupler sends and receives messages to allow the calculation of the surface fluxes. As the interprocessor communications between the atmospheric and oceanic processors are direct, there are no serial bottlenecks. Provided that a reasonably load balanced partitioning of the available processors can be determined, this technique enhances the scalability of the coupled model by increasing overall parallel efficiency.

In this talk, we will review some of our experiences in developing several global models for MPP computers. These include an atmospheric GCM, an oceanic GCM with a dynamic/thermodynamic sea ice model, an ocean biogeochemistry transport model (OBGCM) and an atmospheric chemistry transport model (ACTM). Our current coupled model status will be discussed as well as our strategy for the inclusion of additional component models. Performance issues, particularly load balancing and communication, will also be explored.

*This work was performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract No. W-7405 -Eng-48.